

HOT-SWAPPABLE ROUTER INTERFACE CARD

BACKGROUND OF THE INVENTION

A. Field of the Invention

5 The present invention relates generally to configurable computing devices and, more particularly, to hot-swappable cards that can be inserted and removed from the computing device while the device continues to operate.

B. Description of Related Art

10 Certain computing devices include the ability for the end-user to add and remove hardware components to/from the computing device. For example, a graphics or sound card may be added to a personal computer to increase the performance of the computer. A basic implementation of such a dynamic card insertion requires that the user turn off the computing device before inserting the
15 new card. The computing device then detects, the next time it is turned on, the new hardware card. In this manner, the functionality of the core computing device can be extended by adding an appropriate hardware card.

 Some add-on hardware components are "hot-swappable" (also called "live insertion/removal") meaning that the user does not have to reset the computing
20 device when inserting or removing the hardware component. Instead, the user simply inserts or removes the hardware component while the computing device is operating, and the computing device dynamically recognizes and begins to communicate with the new hardware component. In this manner, the new

component can be inserted and removed from the computing device without interruption of the device.

Hot-swappable add-on components are particularly desirable in critical computing devices that should not be taken off-line. For example, high performance network routers may route traffic through a network for many thousands of users. A router is a device that controls the flow of data packets in a network. It is desirable that these routers operate continuously and without interruption. Accordingly, when inserting add-on components to a high-performance router, it is desirable that the add-on components be hot-swappable.

There are a number of problems that may occur when designing a hot-swappable system for sensitive, high-performance systems. One problem is power disruptions. Inserting a hardware component, such as a circuit board, into an operational system can cause power disruptions due to the sudden increase in load presented to the power supply and the finite response time of that power supply. This is especially true when high capacitance values are used on the circuit board being inserted. This results in an in-rush current which can starve neighboring circuits, potentially causing the logic on these circuits to reset or experience errors.

In addition to power disruptions, software disruptions can occur when inserting and removing a circuit board. In particular, upon removal of a circuit board, the software controlling the system may try to access the board and fail to get a response, thus causing a fatal error. A variation on this problem can occur

on insertion if the software tries to access the circuit board before it is fully inserted and stable in the system.

Thus, there is a need in the art to be able to reliably implement hot-swappable components in a high-performance computing device such as a
5 router.

SUMMARY OF THE INVENTION

Systems and methods consistent with the present invention address this and other need to effectively implement hot-swappable interface cards in a router.

10 In accordance with the purpose of the invention as embodied and broadly described herein, a first aspect of the present invention is directed to a network device including slots for hot-swappable physical interface cards. The network device includes an interface designed to receive the physical interface cards, a plurality of power supply lines, and an on/off power control circuit connected to
15 the plurality of power supply lines. Finally, a controller is connected to the on/off power control circuit. The controller, in response to detecting the presence of the physical interface card in the interface, instructs the on/off power control circuit to turn on the power supply lines by ramping the power supplied to each one of the sequentially turned on power supply lines.

20 A second aspect of the present invention is directed to a hot-swappable physical interface card designed to be inserted into a network device. An interface in the physical interface card connects the physical interface card to the

network device and includes a number of connections. These connections include power supply lines through which power is received from the network device, a high-speed data bus for transmitting and receiving packets of information from the network device, and a control bus. The network device sequentially activates the power supply lines and ramps the power supplied to each one of the sequentially activated power supply lines in response to the network device detecting insertion of the physical interface card.

Additional aspects of the present invention are directed to methods for inserting and removing interface cards from the network device.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an embodiment of the invention and, together with the description, explain the invention. In the drawings,

Fig. 1 is a diagram illustrating functional components of an exemplary network router;

Fig. 2 is a diagram illustrating an implementation of a packet forwarding engine in the router shown in Fig. 1;

Fig. 3 is a block diagram illustrating portions of a flexible PIC controller in the packet forwarding engine shown in Fig. 2;

Fig. 4 is a flow chart illustrating methods consistent with the present invention for inserting a PIC into a flexible PIC controller;

Fig. 5 is a flow chart illustrating methods consistent with the present invention for removing a PIC from a flexible PIC controller; and

Fig. 6 is a diagram illustrating an exemplary graph of a ramped voltage.

5

DETAILED DESCRIPTION

The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents.

10

As described herein, a network router includes hot-swappable physical interface cards that allow the router to communicate using a variety of network technologies. The presence/absence of the interface cards may be detected through both hardware and software. Power to the interface cards may be ramped to avoid disruptive power surges.

15

Fig. 1 is a diagram illustrating, at a high-level, functional components of an exemplary router 100. In general, router 100 receives incoming packets 110, determines the next destination (the next "hop" in the network) for the packets, and outputs the packets as outbound packets 111 on physical links that lead to the next destination. In this manner, packets "hop" from router to router in a network until reaching their final destination.

20

Router 100 includes routing engine 105 and a packet forwarding engine (PFE) 106. Routing engine 105 may maintain one or more routing tables (RTs)

115 and a forwarding table (FT) 116. Through routing tables 115, routing engine 105 consolidates routing information that the routing engine 105 learns from the routing protocols of the network. From this routing information, the routing protocol process may determine the active routes to network destinations and
 5 store these routes in forwarding table 116. Packet forwarding engine 106 may consult forwarding table 116 when determining the next destination for incoming packets 110.

Fig. 2 is a diagram illustrating, in more detail, an implementation of PFE 106 consistent with an aspect of the present invention. PFE 106 generally
 10 handles input/output functions related to the incoming and outgoing packets. More particularly, PFE 106 receives the incoming packets through physical interface cards (PICs) 201, which forward the received packets to flexible PIC concentrators (FPCs) 205. Packet header information from the received packets is transmitted from FPCs 205 to packet processing component 207. In order to
 15 conserve bandwidth in transmitting the header information to packet processing component 207, FPCs 205 may strip the packet header information from the packet body. The body may be stored by FPC 205, with only the header being transmitted to packet processing component 207. The processed header, including its destination PIC, may subsequently be received by FPCs 205, which
 20 reassemble and transmit the packet through the appropriate PIC 201.

Incoming packets 110 and outgoing packets 111 may be received by different PICs 201 over a number of different types of physical transmission standards. For example, the packets may be received through a Gigabit

Ethernet connection, an optical connection (e.g., OC-3 or OC-12) over ATM (asynchronous transfer mode), or an optical connection over SONET (synchronous optical network). Each different physical interface standard may be handled by a specially designed PIC 201. Packets received by the router 100 at one of PICs 201, after processing, may be transmitted from the router 100 by another one of PICs 201. In general, each of PICs 201 perform media-specific details, such as framing and checksum verification, for the packets received by the PIC 201. In this manner, router 100 can support packet routing over a network constructed using heterogeneous transmission components.

FPCs 205 provide a physical interface slot through which each hot-swappable PIC 201 is connected to router 100. In this manner, FPCs 205 connect packet processor 207 to PICs 201. FPCs 205 may perform other packet based functions, such as performing packet integrity checks and applying class-of-service rules to each packet.

Packet processor 207 performs the substantive operations relating to deciding the output port (i.e., the output PIC(s) for each packet) for the incoming packets 110. Packet processor 207 consults forwarding table 116 to obtain this information. In the case of a multicast packet, there may be multiple outgoing interfaces.

As discussed, PICs 201 provide interfaces through which packets can be transmitted over different physical transmission standards. PICs 201 are hot-swappable, giving users the ability to dynamically connect different types of

packet streams to PFE 106 through the appropriate PIC 201, without having to reboot or otherwise disrupt the operation of router 100.

Fig. 3 is a block diagram illustrating portions of FPC 205 consistent with the present invention for implementing a hot-swappable PIC 201. For illustration, two PICs 201 are illustrated as being connected to FPC 205 in Fig. 3. In practice, each FPC 205 may support more than two PICs 201 (such as the four PICs illustrated in Fig. 2).

FPCs 205 may connect to hot-swappable PICs 201 through a physical interface 310. Through each physical interface 310, a number of different data and power signals are transmitted to the corresponding PIC 201. In particular, as shown, FPC 205 includes a 1.5 volt power source 315, a 2.5 volt power source 316, a 3.3 volt power source power source 317, and a 5 volt power source 318. The 1.5 volt source 315 may connect to PIC 201 through switch 320 and power line 321. Similarly, the 2.5 volt source 316 may connect to PIC 201 through switch 322 and power line 323; the 3.3 volt source 317 may connect to PIC 201 through switch 324 and power line 325; and the 5 volt source 318 may connect to PIC 201 through switch 326 and power line 327. Data buses 330 and 331 connect PIC 201 to FPC controller 335 and packet I/O manager 336, respectively. Data bus 331 may include a high speed data bus over which the packets received at PIC 201 are transmitted. Data bus 330 may include a pin detection line 340 and a control bus 341. Control bus 341 may be, for example, a parallel CPU bus. FPC 205 additionally includes for each PIC 201 connected to the FPC 205, circuitry for managing power supplied to the PIC, illustrated as

on/off control circuit 338 and power monitor circuit 339. A serial bus controller 352, positioned between FPC 335 and on/off control circuit 338, may be used to facilitate communications between FPC controller 335 and on/off control circuit 338.

5 Packet I/O manager 336 receives packets from PICs 201 over high-speed data bus 331. Packet I/O manager 336 may transmit the packet header information to packet processing component 207 and store the packet body in memory 337. Memory 337, in addition to being used by packet I/O manager 336, may be shared with packet I/O managers on other FPCs 205. In this manner,
10 the collection of memories 337 on all of the FPCs 205 can together function as one larger distributed memory.

As previously mentioned, FPC controller 335 generally manages the interfacing of PICs 201 to FPC 205. More particularly, status and control information relating to PIC 201 may be exchanged on control bus 341 and the
15 insertion state of PIC 201 may be communicated using pin detection line 340. A more detailed description of the operation of the components shown in Fig. 3 for implementing hot-swappable PICs 201 will be further discussed with reference to Figs. 4 through 6.

Fig. 4 is a flow chart illustrating methods consistent with the present
20 invention for inserting one of PICs 201 into an FPC 205.

FPC controller 335 begins a PIC insertion procedure when it detects the presence of a PIC 201 in physical interface 310. (Act 401). The presence of PIC 201 may be initially detected through an electrical connection with a

predetermined pin in the PIC 201 connected to pin detection line 340. The pin may be a "short" pin that implements a "Last-Make First-Break" connection with physical interface 310. More specifically, the predetermined pin may be designed to be the last pin to come into contact with physical interface 310 when the PIC is inserted and the first pin to lose contact with physical interface 310 when the PIC is removed. Accordingly, when the presence of PIC 201 is detected on pin detection line 340, FPC controller 335 can be sure that the PIC 210 is fully inserted. A filter may be inserted between physical interface 310 and FPC controller 335 to eliminate signal bounce.

After detecting the presence of the newly inserted PIC 201, FPC controller 335 supplies power to the detected PIC 201. More particularly, the FPC controller signals the on/off control circuit to begin supplying power to the power lines 321, 323, 325, 327. (Act 402). On/off control circuit 338 subsequently ramps the power supplied to each power line. (Act 403). Monitor circuit 339 monitors the newly supplied power to ensure stability of the power supply. (Act 404). If the power has not stabilized after a predetermined time period, FPC controller 335 restarts the power-up procedure.

On/off circuit 338, instead of abruptly turning the power on or off through its associated switches 320, 322, 324, and 326, may activate the power in a ramp fashion to thus produce a corresponding ramp in the applied or removed power. Fig. 6 is a diagram illustrating exemplary ramped voltages. As shown, the supplied voltage is steadily increased to the load PIC over time. A typical ramping time may be, in the range of, 5 milliseconds to 20 milliseconds.

Similarly, on power-down, the supplied power may be ramped down. By ramping the supplied voltage, a more stable on/off procedure can be achieved.

Fig. 5 is a flow chart illustrating methods consistent with the present invention for removing a PIC from an FPC.

5 The FPC controller 335 may begin by detecting the removal of the PIC 201. (Act 501). FPC 205 detects the removal of PIC 201 by the break of the electrical connection indicated on detection pin line 340. In addition, or in the alternative, FPC controller 335 can detect the removal of PIC 201 in software by detecting a time-out on control bus 341 after the PIC 201 fails to respond to bus
10 requests for a predetermined period of time (e.g., 3 microseconds).

Dynamically removing one of PICs 201 can create spurious signals on control bus 341 and high-speed data bus 331. These signals can confuse or even damage packet I/O manager 336 and/or FPC controller 335. Accordingly, FPC controller 335, after detecting the removal of PIC 201, blocks the reception
15 of any control signals from PIC 201 over control bus 341 and zeroes out data lines on bus 341. (Act 502). FPC controller 335 may additionally assert a signal to packet I/O manager 336 to clear the data lines on high-speed bus 331. (Act 503). Finally, the FPC controller 335 turns-off the power supplied through power lines 321, 323, 325, and 327. (Act 504). As with the power-up sequence, FPC
20 controller 335 may power-down the supplied power through on/off control circuit 338 by ramping down each power supply (as shown in Fig. 6).

Acts 501-504, as performed by FPC controller 335, allow PICs 201 to be removed from router 100 while minimizing the chance of experiencing errors or

electrical disruption of data or control signals on FPCs 205. Further, by sequentially ramping-up and ramping-down the power supplied to PICs 201, undesirable power surges caused by capacitive coupling can be reduced.

5 Buffers, such as buffers 350 and 351, may be inserted between PICs 201 and the FPC controller 335 and/or packet I/O manager 336 to further reduce the occurrence of spurious signals.

Although FPC controller 335 is shown as a single controller for FPC 205, in an alternate embodiment, FPC controller 335 may be implemented as independent controllers for each PIC 201. For example, for each PIC 201,
10 controller 335 may include a separate CPU used to control the PIC 201. This can be desirable to insure that each PIC 201 operates independently of the other PICs and that an error in one PIC does affect other PICs.

As described above, hot-swappable physical interface cards allow a router to communicate using a variety of network technologies. The physical interface
15 cards are not required to include any internal power supplies and can be freely exchanged with each other.

The foregoing description of preferred embodiments of the present invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications
20 and variations are possible in light of the above teachings or may be acquired from practice of the invention.

Although described as being primarily implemented in hardware, the present invention may be embodied in hardware and/or in software (including

- firmware, resident software, micro-code, etc.). Furthermore, the present invention may take the form of a computer program product on a computer-usable or computer-readable storage medium having computer-usable or computer-readable program code embodied in the medium for use by or in connection with
- 5 an instruction execution system.

The scope of the invention is defined by the claims and their equivalents.

13